

Importance-Driven Rendering in Interventional Imaging

Martin Haidacher
martin.haidacher@gmail.com

Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria
Imaging & Visualization Department, Siemens Corporate Research, Princeton, USA




Figure 1: The images (b) to (d) show step by step the influence of all three modifications of optical properties by a single importance value compared to standard volume rendering in (a).

Abstract

In this work a combined visualization of dense clinical data like 3D CT (Computed Tomography) together with real-time images of medical intervention applications is presented. The main challenge here is to provide a fused visualization that allows sufficient spatial perception of the important parts, as derived from the pre-operative data, while not occluding the information in the real-time image embedded within the volume.

This paper presents a new approach of importance definition for volumetric data and how this importance can be used to create a feature emphasized visualization. Furthermore the viewpoint and the position of the intervention image is used to generate a contextual cutaway which influences the density of the visualization to avoid an occlusion of the real-time image by less important parts of the volumetric data.

Finally this new approach is used to show results in a clinical context where a pre-operative CT scan is visualized alongside a tracked ultrasound image, such that the important vasculature is depicted between the viewpoint and the ultrasound image, while a more opaque representation of the anatomy is exposed in the surrounding area.

Keywords: Importance-Driven Rendering, Interventional Imaging, Contextual Cutaway, Transfer Function, Ultrasound, CTA, CT, MRI, Liver Biopsy

1 Introduction

Medical visualization is an active research area since many years. The development of more accurate imaging systems for medical examination makes it necessary to use an adapted visualization to handle these large data sets. Therefore such a visualization need to emphasize the important parts of the whole data set for a specific medical task. The less important parts are only used as contextual background information. The reduction of the information to the important parts makes it faster, saver, and more reliable for a physician to manage a particular task. Depending on the imaging system and on the medical task, different approaches have been developed in the last years to reach this goal. This paper presents a new technique which can be especially used for interventional imaging in combination with co-registered pre-operative data sets.

Interventional imaging is a very often used technique for less immersive operations. A real-time imaging system is used during the intervention. In this work the main focus lies on the use of ultrasound as the imaging system. A liver biopsy, e.g., is a common task where an ultrasound imaging system is being used to guide the needle to the target position in the liver. The critical point in this intervention is to navigate the needle through the liver without injuring one of the blood vessels in the liver. To avoid such an injury an additional pre-operative CTA (Computed Tomography Angiography) or CT scan is used to plan the

needle path before the intervention.

This paper shows how to use this pre-operative data merged with the ultrasound image directly in the intervention. The inherent problem here is that it is difficult to show enough data from the dense pre-operative scan without occluding the ultrasound image plane. Motivated by this problem, a new approach for visualizing an object of interest embedded within volumetric data is presented. This approach ranks and emphasizes materials within the volume by a user-specified importance and adaptively changes the amount of information displayed from the volume based on the viewing direction and the location of the ultrasound plane.

This approach can be logically split into four components, as described in Section 3. The first is the assignment of importance values to materials within the volume which is done by adding an extra component to the transfer function specification (3.1). Next, the rendered materials are emphasized by their assigned importance, by modifying optical properties (3.2). Finally, a flexible, view-dependent cutaway structure is defined (3.3) and this definition is used to cut away occluding material based on its importance (3.4). This results in a visualization where the object of interest embedded in the volumetric data is clearly visible together with parts of the volumetric data deemed especially important while progressively trimming away material of lower importance. In Section 4 concrete medical applications for the use of this new approach are presented.

2 Related Work

This work aims at one of the main challenges in visualization research to satisfy needs of a specific diagnostic and treatment environment. The challenge is to provide clear understanding of the complex data and guide the user to the most relevant information. The relevance of features in the data can be assigned in many different ways, i.e., from user steered segmentation to a fully automatic process. Relevance then serves as the controlling parameter for assigning visual representations among those features. The combination of dense (visually prominent) visual representation of most relevant features with sparse (visually suppressed) information about other features is often denoted as focus+context visualization [8]. Focus, i.e. the most relevant feature, is represented very densely and context is presented sparsely to indicate overview information.

In the visualization of 3D flow data a degree of interest function (DOI) has been introduced to include the user interest in the visualization [9]. The DOI then affects previously defined optical properties to modulate the opacity or change the color hue or saturation. In the volume visualization of scalar data two-level volume rendering introduced the notion of objects in the data where each object can be represented with a different rendering technique [10].

In the above mentioned cases the viewpoint position is not considered. The assignment of a dense visual representation to the most important focus region does not automatically guarantee a clear view at the focus. To achieve this, optical properties of a context region that occludes the focus have to be adapted. Several smart visibility techniques have been developed for this purpose, often inspired by traditional illustration techniques [17].

One of the approaches, which is using the data relevance to automatically generate cut-away visualizations, is importance-driven visualization [18]. Here several view-dependent operators have been introduced to suppress occluding context areas in order to increase the visual prominence of the focus region. In contrast to this work, the presented approach does not require segmentation information, as the importance classification is done in the data value space. Similarly to the presented approach context-preserving volume rendering [1] or opacity peeling [14] also do not use any segmentation information and generate special kind of cut-away views. These concepts, however, do not incorporate an importance assignment mechanism so the object emphasis and suppression mechanism is harder to specify.

Occlusion of most important information is currently a very active research area. Many interactive techniques for cut-away (or ghosted) [2, 13], peel-away [5] and exploded visualizations [3] have been developed. Smart visibility techniques have been applied in the visualization of the oil and gas data [16]. In the domain of medical imaging, cut-away visualizations have been applied for the visualization of peripheral arteries in lower extremities [15]. Another example from the medical imaging domain is neck dissection planning for enlarged lymph nodes removal [12].

There are two major challenges associated with multimodal medical data. The first challenge is the data registration. For this work the registration of the CT (or MRI) data with the ultrasound data is being done as described in the previous work [19]. The second challenge is how to combine multiple modalities in a single visualization. Magic mirrors have been applied for multimodal visualization of the human brain [11]. Another way of multimodal visualization is to combine multiple modalities or properties through smooth brushing and linking of scatter plots [6]. The product of this multimodal selection is mapped to one-dimensional DOI classification and DOI is directly mapped to visual representation. Multiple modalities are here used for selection of a *feature*. According to the needs of the interventional setup, importance parameterization is used as a steering mechanism for combination of data from different modalities. This means that multiple modalities need to be explicitly present.

This type of visualization has been firstly released in our previous paper [4].

3 Occlusion-Driven Feature Emphasis




Figure 2: Overview of the pipeline feature emphasis and contextual cutaway visualization.

This section describes all steps of the pipeline for importance-driven rendering in interventional imaging. Figure 2 shows an overview of this pipeline. The top three nodes are the input sources and the subsequent nodes are processing steps described in the following subsections.

The base inputs are the volumetric data and the object of interest, which occupies a region in the volumetric data set. In this case a CTA data set is used as volumetric data and an ultrasound plane is the object of interest. The main goal of the approach is to have a combined visualization of the volumetric data with the ultrasound plane while the ultrasound plane should not be occluded by less important parts of the volumetric data set.

In the left pipeline, the standard color and opacity transfer function is augmented with an importance value which specifies the relative importance of each regional component of the transfer function (3.1). In Section 3.2 this importance value is used to modify some optical properties to emphasize more important parts.

In the right pipeline, the object of interest is used with the viewpoint information to create the contextual cutaway structure, which divides space into different contextual regions and defines a function for determining occlusion (3.3).

Finally the occlusion for a given sample point is used to modify the importance value and furthermore the optical properties. In Section 3.4 the formulas for modifying the importance and the visual effect on the final appearance of a rendered sample point are described.

3.1 Defining Importance via Transfer Function

Due to the mass of information in a single volume data set a representation is needed which directs the observer to the most significant parts of the whole data set. For that reason a kind of segmentation is needed which splits the

volume into more and less important parts. The used visualization technique should then be able to highlight the more important parts in addition with an overview visualization of the less important parts. The overall goal is to make the cognition of important parts fast and easy while keeping an impression of the position and the relation of these parts with respect to the rest of the data set.

Importance-driven rendering, introduced by Viola et al. [18], is one approach to reach these goals. As input a data set is needed with a segmentation of the different parts according to their spatial location. To segment a volume data set in spatial areas is a time consuming and computationally intensive task which mostly also needs some user interaction. This approach avoids this additional step to reduce the pre-computation time and to reduce user interaction. The last issue is very important because for a practical use of the method it is not desired by a physician to, e.g. spend some time in selecting seed points for a segmentation.

With this background a new importance-driven rendering technique is introduced which segments the different important parts not by their spatial position but only according to their position in the transfer function space. This has the benefit that a segmentation can be pre-defined once for a given kind of medical task on a given machine. It works especially good for data sets which are acquired with the additional use of a contrast agent.

The original idea of a transfer function is the assignment of a color and opacity to every sample point in the volume data set. For this reason a feature vector \vec{x} is generated from the local properties of a single sample point.

$$g(\vec{x}) \rightarrow (\vec{c}, \alpha)$$

The tuple (\vec{c}, α) defines the RGB color as vector \vec{c} and the opacity (α) for the given feature vector. The number of elements for the feature vector defines the complexity of the transfer function. Mostly the intensity value is used as only element for the classification. Additionally also the gradient magnitude, curvature or other higher order properties can be used for the classification. The transfer function can be split into different components which are defined by their color and opacity distribution in the transfer function space. Each of this components should belong to a tissue in the volumetric data set. The idea behind using importance values for the transfer function is to assign an additional single value to each component of the transfer function. For a single point in the transfer function space two or more components can overlap each other. In such a case a blending of the color values, opacity, and importance is necessary. The following formula describes this blending:

$$g(\vec{x}) \rightarrow \left(\frac{\sum_{i=1}^{N_{comp}} \vec{c}_i * \alpha_i * I_i}{\sum_{i=1}^{N_{comp}} \alpha_i * I_i}, \max_{i=1..N_{comp}} \alpha_i, \max_{i=1..N_{comp}} I_i \right)$$

I_i are the importance values for each component i . N_{comp} is the number of all components overlapping at a single

point. The RGB color components are combined in this formula in the vector \vec{c} . This blending is only done once before the rendering and results in a lookup-table which assigns each point in the transfer function space to a quintuple consisting of RGB color values, an opacity value, and an importance value.

3.2 Importance-Driven Emphasizing

The importance value introduced in the previous section can be used to modify the optical properties of a sample point in the volume, with the goal of emphasizing important parts. Dependent on this importance value the properties should be modified. The following properties are chosen to be modified in the following order:

- Shading
- Opacity
- Silhouette Enhancement

The goal of emphasis is to guide the viewer quickly to the most important parts of an image, which can be achieved by rendering more details for more important materials. The less important materials should not occlude more important parts. This can be done by modifying each of the mentioned properties by the importance value.

The first property, the shading, should be more realistic for important parts. For that reason no shading is used for unimportant parts and a more complex shading method for important parts. The following formula describes the modification of the shading by the importance value.

$$\vec{c} = \vec{c}_{shaded} * I + \vec{c}_{unshaded} * (1 - I)$$

\vec{c}_{shaded} is the resulting color from the more complex shading method and $\vec{c}_{unshaded}$ the color resulting from no shading. The importance value I is used as interpolation value between these colors. Figure 1 (b) shows the result of this modification in contrast to standard volume rendering without using an importance value in Figure 1 (a). The vessels are most important in this case.

As one can see this shading modification makes less important parts more opaque than with normal shading. Reducing the opacity by the importance value can eliminate this penalty. To keep the outline of the less important parts this reduction is done dependent on the dot-product of the normal vector \vec{n} and the viewing direction \vec{v} .

$$\alpha = \alpha * \max\left(I, (1 - |\vec{n} \cdot \vec{v}|)^\delta\right) * I^{\frac{1}{\delta}}$$

The \max function is needed to reduce the opacity just for small importance values I and for sample points with a normal vector almost parallel to the viewing vector. The δ exponent is used to control the opacity reduction. A higher value makes the less important parts more transparent. The additional multiplication with the importance value reduces the opacity in general for less important




Figure 3: Comparison between importance-driven Phong shading (a) and Gooch shading. The spatial perceptibility is a little bit better with Gooch shading.

parts. Figure 1 (c) shows the result of the additional opacity modification together with the shading modification.

The silhouette enhancement should highlight the outline of important parts to get a higher contrast for these parts. To achieve this the color at the silhouette outline of more important parts is darkened. The modification is dependent on the importance value and on the dot-product between the viewing vector and the normal vector. The darkening is highest for important parts and sample points with a normal vector which is perpendicular to the viewing direction.

$$\vec{c} = \vec{c} * \max(1 - I, (|\vec{n} \cdot \vec{v}|)^\epsilon)$$

The exponent ϵ effects the thickness of the outline. The smaller this value the thicker the outline. Figure 1 (d) shows the final rendering result of all three modifications. Compared with Figure 1 (a) the effect of the modifications is clearly visible. Even the vessels behind the skin and bones can be seen now without using any clipping.

To get more contrast for the important parts Phong shading was compared with other shading techniques. Using Gooch cool-to-warm shading [7] gives a stronger contrast between important and non-important materials, as well as increased spatial perceptibility. In Figure 3 you can see a comparison between Phong and Gooch shading.

Until now a method to emphasize important parts in a volumetric data set was introduced. The following sections will now concentrate on the combination of this technique with a region of interest embedded in the volume to get a fused visualization.

3.3 Contextual Cutaway Views

As stated earlier, our goal is to present the object of interest in such a way that the material around it is cut away based on the importance of the material. In this section a flexible cutaway structure is presented that ultimately allows us to trim away materials of different importance at different levels of the cutaway structure.

The simple cutaway view definition partitions space into two distinct regions: the area inside the cutaway, which is denoted as *clear* region, and everything else, which is denoted as *base*. This is formalized by defining an *occlusion function*, denoted Ω . This occlusion function will represent the degree to which a point in space occludes the object of interest. In the simple cutaway scenario, at a given point, $\Omega = 1$ if the point is inside the *clear* region and 0 if it is inside the *base* region. In eye space, a cutaway surface can be represented by a depth function $\xi(\theta) = z$, where z is the depth of the cutaway surface with angle θ at a given point projected onto the surface. We can then define Ω for a given point in eye space as follows, where p_z is the z component of the point and $\text{step}(a, x) = 0$ if $x < a$ and 1 if $x \geq a$.

$$\Omega = \text{step}(\xi(\theta), p_z)$$

This binary definition suggests rendering can have only two modes: sparse (for the *clear* region) and dense (for the *base* region). In order to give more control over the rendering of materials with multiple importance values, a new cutaway definition is proposed where occlusion values vary smoothly between 0 and 1 over 3D space.

As beginning of the modification of the simple cutaway definition a second cutaway surface is included defined by a wider angle. This new region, which is denoted as *transition* region, can have an occlusion function that varies between the two cutout surfaces. This is defined as shown in Figure 4, allowing us to determine the cutout angle of points located in the *transition* region, relative to the two bounding angles. This will ultimately allow variation of visibility in the image outside the projected object-of-interest silhouette by letting us cut or fade away materials at varying angles.

To control the visibility in the image over the object-of-interest, another region is added, the *overlay* region. This region is bounded by the cutaway surface of θ_1 offset a thickness d towards the camera, as shown in Figure 4.

Considering these four regions, the occlusion function for a given point in eye space is defined as follows, where θ_1 and θ_2 are the cutaway angles, d is the thickness of the *overlay* region, and $\text{ramp}(a, b, x) = 0$ if $x \leq a$, 1 if $x \geq b$, and is a linear ramp from 0 to 1 for $a < x \leq b$.

$$\Omega = \frac{\text{ramp}(\xi(\theta_2), \xi(\theta_1), p_z) + \text{ramp}(\xi(\theta_1), \xi(\theta_1) + d, p_z)}{2}$$

This definition results in $\Omega = 0$ for points in the *base* area, $\Omega = 0.5$ for points on the *transition-overlay* boundary, and $\Omega = 1$ for points in the *clear* area, with the appropriate linear blends for points in between the boundaries, as shown in Figure 4.

3.4 Modifying Importance by Occlusion

In the previous section a flexible cutaway structure was defined which can be combined now with the rendering technique of Section 3.2. The occlusion function can be




Figure 4: Given an object of interest (shaded) and a viewing direction, a cutaway with distinct *base*, *transition*, *overlay*, and *clear* regions can be constructed, as shown in this cross section diagram. Then an occlusion function Ω in 3D space is defined to modify opacity in different areas of the cutaway structure. In the traditional simple cutaway, $\theta_1 = \theta_2$ and $d = 0$.

used to modify the importance value for different regions in the cutaway area. For the *base* area we want to have a very dense visualization to get a good impression where the cutaway area starts. The *transition* area should show a smooth fading between the dense visualization and an emphasis visualization of the important parts like in Section 3.2. In the *overlay* area only the important parts should be visible. Finally in the *clear* area everything should be cut away. The modification of the importance value is done by the following formula:

$$I_m = \text{ramp}(\tau_l, \tau_u, I)$$

The *ramp* function is the same function as described in the






Figure 6: Relation between Ω and the importance modification.


section before. I is the original importance value from the




(a) $\theta_1 = \theta_2, d = 0$



(b) $\theta_1 < \theta_2, d = 0$



(c) $\theta_1 < \theta_2, d > 0, \sigma = 0.1$



(d) $\theta_1 < \theta_2, d > 0, \sigma = 10$

Figure 5: The effects of various parameters in the cutaway structure are shown in the above images. Image (a) shows a traditional cutaway view where the other three images show a smooth transition and an overlay area in front of the important part.

transfer function and I_m is the modified importance value. τ_l and τ_u are the threshold values for the ramp which are dependent on the Ω from the contextual cutaway.


$$\begin{aligned} \tau_l &= \left(\frac{\text{floor}(2 * \Omega)}{2} \right)^\sigma \\ \tau_u &= \min(1, 2 * \Omega) \end{aligned}$$

Figure 6 shows the relation between Ω and the modified importance. You can see that the lower threshold of the ramp τ_l is always smaller or equal to the upper threshold τ_u . For all doubles (Ω, I) which are above the τ_u line in the diagram the modified importance will be 1. If a sample point from the volume lies on or below the τ_l line then I_m will be 0. For all points between these two lines the modified importance is calculated by a linear interpolation. For an Ω value of 0 in the *base* area I_m will be 1 for all importance values greater than 0. This results in a dense visualization of all parts.


In the *transition* area the lower threshold is always 0 while the upper threshold increases to 1 towards the border to the *overlay* area. This increase of τ_u fades out the

less important parts. In the *overlay* area τ_u stays constant at 1 and τ_l jumps from 0 in the *transition* area to 0.5^σ . The σ is used to modify the lower threshold value for this area. A high σ makes the threshold lower and a low σ makes it higher. The lower this threshold is the more parts with a lower importance can be seen in this *overlay* area. Finally in the *clear* area with $\Omega = 1$ all importance values are modified to 0 because both threshold values are at 1 and therefore everything is clipped away in this area.

In Figure 5 you can see examples with different parameters for the cutaway angles, the *overlay* thickness, and the τ_l exponent σ . Figure 5 (a) shows an example of a traditional cutaway view. The *transition* area and the *overlay* area are not present. In Figure 5 (b) the angles Θ_1 and Θ_2 are different and therefore the *transition* area is not zero. A smooth fade out of less important parts can be seen. In Figure 5 (c) and (d) the thickness of the *overlay* area is set to a value greater than 0. Additionally a higher σ value is used in Figure 5 (d) so the lower threshold in the *overlay* area is lower and more less important parts like the hull of the lung can be seen in this area.



(a) Longitudinal ultrasound image, small d , Phong shading



(b) Longitudinal ultrasound image, large d , Gooch shading

Figure 7: Ultrasound images of a liver in the context of a CTA scan.

4 Results

For the results ultrasound exams recorded on patients and volunteers are used, where the spatial location and orientation of the transducer has been tracked using a magnetic position sensing system. The ultrasound sweeps have been manually registered (i.e. spatially aligned) to the tomographic 3D scans, aided by automatic techniques based on previous work [19]. To show interventional ultrasound in the context of organ vasculature, an early arterial phase CTA scan from a patient's liver is used. The injected contrast agent causes the vascularity to show up with high intensities in the scan, easily distinguishable from surrounding tissue. Figure 7(a) depicts the proposed rendering for a longitudinal ultrasound image of the liver and the corresponding CTA data. In front of the ultrasound only the vessels and parts of the skin are rendered. The importance value of the vessels is 1 while the importance value of the skin is 0.3. This is the reason why the skin is rendered more transparent and with less details. All other tissues have an importance value lower than the τ_l threshold, so they are not visible in the *overlay* area. The thickness d is set to a small value to just show the vessels in a near distance to the ultrasound plane. Figure 7(b) shows a similar rendering with Gooch shading instead of Phong shading. The thickness is set to a higher value to have a larger *overlay* area. In Figure 8 a similar visualization is applied using a regular CT scan without contrast agent. Nevertheless a physician get a good impression of the position of the ultrasound plane in relation to the body. With the magnetic tracking and the automatic registration this technique is a very good tool to explore the body by using the ultrasound probe as interaction device.

In all data sets, the ultrasound images are made clearly visible and their spatial context within the 3D volume is

well defined due to the view dependent cutaway structure. In addition to the global 3D relation, critical anatomical structures such as liver vasculature that must not be punctured in an interventional scenario, can be visualized as well. This is important for image-guided needle procedures and can increase the acceptance of multimodal visualization for such procedures, as the physician always has a concise view of all the needed information, regardless of the ultrasound probe position.

Finally, the presented techniques have been designed for and tested with a hardware-based GPU raycaster. As such, the system can operate at interactive frame rates, which is necessary for use in a clinical scenario.

5 Conclusion

The paper presented a suite of new techniques for visualizing objects of interest embedded within volumetric data sets. These techniques allow for emphasis of important volumetric features as defined by transfer functions as well as visualization of contextual information relative to the object of interest by means of a new, flexible cutaway structure. The techniques were applied to the area of image-guided needle procedures, and show fused visualizations of ultrasound with CTA and CT volumes.

Contemporary radiology software provides a built-in set of transfer functions adapted to various imaging modalities, which doctors use to visualize patient data and distinguish various tissues and organs. As the introduced algorithm classifies importance of materials in the transfer function space, it would be reasonable to expect that these transfer functions could be augmented with importance information for various medical scenarios. As such, the presented techniques can have a positive impact on the medical visualization community.




Figure 8: Transversal ultrasound image of a patient’s liver, in the context of a (non-contrasted) CT scan, precisely aligned. Bones ($I = 1$), skin ($I = 0.3$), and colon ($I = 0.3$) have a higher importance than the rest.

References

- [1] S. Bruckner, S. Grimm, A. Kanitsar, and M. E. Gröller. Illustrative context-preserving exploration of volume data. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1559–1569, 2006.
- [2] S. Bruckner and M. E. Gröller. VolumeShop: An interactive system for direct volume illustration. In *Proceedings of IEEE Visualization’05*, pages 671–678, 2005. <http://www.volumeshop.org/>.
- [3] S. Bruckner and M. E. Gröller. Exploded views for volume data. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1077–1084, 2006.
- [4] M. Burns, M. Haidacher, W. Wein, I. Viola, and M. E. Gröller. Feature emphasis and contextual cutaways for multimodal medical visualization. (accepted for EuroVis 2007).
- [5] C. Correa, D. Silver, and M. Chen. Feature aligned volume manipulation for illustration and visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1069–1076, 2006.
- [6] H. Doleisch, M. Gasser, and H. Hauser. Interactive feature specification for focus+context visualization of complex simulation. In *Proceedings of VisSym’03*, pages 239–248, 2003.
- [7] A. Gooch, B. Gooch, P. Shirley, and E. Cohen. A non-photorealistic lighting model for automatic technical illustration. In *SIGGRAPH ’98: Proceedings of the 25th annual conference on computer graphics and interactive techniques*, pages 447–452, 1998.
- [8] H. Hauser. *Scientific Visualization: The Visual Extraction of Knowledge from Data*, chapter Generalizing Focus+Context Visualization, pages 305–327. Springer, 2005.
- [9] H. Hauser and M. Mlejnek. Interactive volume visualization of complex flow semantics. In *Proceedings of VMV ’03*, pages 191–198, 2003.
- [10] H. Hauser, L. Mroz, G. I. Bischi, and M. E. Gröller. Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):242–252, 2001.
- [11] A. König, H. Doleisch, and M. E. Gröller. Multiple views and magic mirrors - fMRI visualization of the human brain. In *Proceedings of SCCG ’99*, pages 130–139, 1999.
- [12] A. Krüger, C. Tietjen, J. Hintze, B. Preim, I. Herpel, and G. Strauß. Interactive visualization for neck dissection planning. In *Proceedings of EuroVis’05*, pages 295–302, 2005.
- [13] J. Krüger, J. Schneider, and R. Westermann. Clearview: An interactive context preserving hotspot visualization technique. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):941–948, 2006.
- [14] C. Rezk-Salama and A. Kolb. Opacity peeling for direct volume rendering. *Computer Graphics Forum*, 25(3):597–606, 2006.
- [15] M. Straka, M. Červeňanský, A. La Cruz, A. Köchl, M. Šrámek, M. E. Gröller, and D. Fleischmann. The VesselGlyph: Focus & context visualization in CT-angiography. In *Proceedings of IEEE Visualization’04*, pages 385–392, 2004.
- [16] K. Hinrichs T. Ropinski, F. Steinicke. Visual exploration of seismic volume datasets. In *Proceedings of the WSCG ’06*, pages 73–80, 2006.
- [17] I. Viola and M. E. Gröller. Smart visibility in visualization. In *Proceedings of Computational Aesthetics in Graphics, Visualization and Imaging*, pages 209–216, 2005.
- [18] I. Viola, A. Kanitsar, and M. E. Gröller. Importance-driven feature enhancement in volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):408–418, 2005.
- [19] W. Wein, B. Röper, and N. Navab. Automatic registration and fusion of ultrasound with CT for radiotherapy. In *Proceedings of MICCAI ’05, Lecture Notes in Computer Science*, pages 303–311. Springer, 2005.